GFlowNets for Learning Better Drug-Drug Interaction Representations

Azmine Toushik Wasi
[1]Computational Intelligence and Operations Laboratory (CIOL)
[2]Shahjalal University of Science and Technology
azmine32@student.sust.edu

Abstract

Drug–drug interactions pose a significant challenge in clinical pharmacology, with severe class imbalance among interaction types limiting the effectiveness of predictive models. Common interactions dominate datasets, while rare but critical interactions remain underrepresented, leading to poor model performance on infrequent cases. Existing methods often treat DDI prediction as a binary problem, ignoring class-specific nuances and exacerbating bias toward frequent interactions. To address this, we propose a framework combining Generative Flow Networks (GFlowNet) with Variational Graph Autoencoders (VGAE) to generate synthetic samples for rare classes, improving model balance and generating effective and novel DDI pairs. Our approach enhances predictive performance across interaction types, ensuring better clinical reliability.

1 Introduction

Drug–drug interactions (DDIs) represent a critical issue in clinical pharmacology, as adverse interactions can lead to significant patient harm and reduced therapeutic efficacy [27]. Numerous computational models have been developed to predict DDIs, leveraging diverse features from chemical structures to biological networks [3, 26]. However, a pervasive challenge in this domain is the severe class imbalance among interaction types. Common interaction types, such as synergistic effects or well-characterized adverse reactions, dominate the datasets, while rare interaction types remain under-represented [5]. This imbalance hinders the model's ability to learn nuanced patterns associated with infrequent interactions. As a consequence, predictive performance on rare, yet clinically significant, interaction types suffers. The disparity in data distribution thus poses an urgent need for innovative solutions in DDI prediction.

In light of the class imbalance, existing state-of-the-art methods are often trained in a binary setting [3, 23, 14], treating the DDI prediction task as a simple presence-or-absence problem. This binary framing tends to disregard the inherent heterogeneity among different interaction types. Consequently, the models become biased towards common interaction types and fail to adequately capture the underlying characteristics of rare classes. The motivation for this work stems from the recognition that a one-size-fits-all approach is insufficient for capturing the diversity of DDIs. Addressing the imbalance is crucial for improving the reliability and clinical utility of these predictions [3]. By acknowledging and explicitly targeting the disparity in interaction frequencies, we aim to bridge the gap between theoretical performance and real-world application with GFlowNets [16, 18].

To mitigate the challenges posed by class imbalance, we propose an innovative framework that integrates a Generative Flow Network (GFlowNet) [1, 8] module with a Variational Graph Autoencoder (VGAE) [14, 23]. Our approach first computes a reward for each interaction type that is inversely

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: NeurIPS 2025 Workshop on Structured Probabilistic Inference & Generative Modeling: Probabilistic Inference in the Era of Large Foundation Models.

proportional to its frequency, thereby guiding the sampling process towards under-represented classes. The GFlowNet module sequentially generates synthetic DDI samples by first selecting an interaction type based on this reward and then sampling a drug pair conditioned on that type. These synthetic samples are then used to augment the original training data, effectively balancing the class distribution. Experimental results indicate that this method enhances the model's ability to predict both common and rare interaction types with improved robustness. The proposed framework not only addresses a critical limitation in current DDI prediction models but also holds promise for broader applications in imbalanced classification problems across biomedical domains.

2 Related Works

2.1 Generative Flow Networks (GFlowNets)

Generative Flow Networks (GFlowNets) [1, 8] are emerging as a powerful paradigm in drug discovery by enabling efficient exploration of the vast chemical space. In structure-based drug design (SBDD), they incorporate geometric information from protein–ligand interactions, leading to improved binding affinities and outperforming conventional approaches through the use of trigonometrically consistent embeddings [**?**]. Beyond binding, GFlowNets also address synthesizability, as demonstrated by the RxnFlow model, which integrates molecular building blocks and reaction templates to achieve a synthesizability rate of 34.8% while excelling in pocket-specific optimizations [19]. To further expand chemical space exploration, Atomic GFlowNets (A-GFNs) generate molecules atom by atom, guided by proxy rewards from inexpensive molecular descriptors that align the search with desirable pharmacological properties [17]. These advances position GFlowNets as versatile tools that combine structural accuracy, chemical feasibility, and broad diversity in molecular generation. However, challenges persist in ensuring interpretability and explainability, as the black-box nature of these models can hinder trust and limit adoption in high-stakes biomedical applications [11]. Overall, GFlowNets represent a transformative shift in computational drug discovery, bridging structure, synthesis, and pharmacological relevance in a unified generative framework.

2.2 Drug Drug Interaction

Drug discovery tasks like drug combination and drug-drug interaction (DDI) prediction have been widely studied. DDI prediction, the focus of this work, aims to identify synergistic, antagonistic, or neutral interactions between drugs. Recent advances leverage graph-based architectures like MFConv [4] and GraphDTA [15] for molecular representation. Lim et al. [12] integrate 3D structures and graph attention, while Torng and Altman [21] employ graph autoencoders for protein-pocket representations. Autoencoders such as VGAE [10] and GraphVAE [20] generate latent graph embeddings, aiding in molecular structure modeling. GraphMAE [7] introduces masked pretraining for robust graph learning. Knowledge-based autoencoders like [3] incorporate semantic embeddings for DDI tasks. Self-supervised learning methods such as SupDTI [2] and SuperGAT [9] improve node representations using unlabeled data. Contrastive frameworks like HeCo [22] and SHGP [25] enhance graph pretraining via meta-paths and structure-aware pseudo-labeling. Our work builds on these ideas by proposing a VAE-based DDI prediction model enriched with graph-aware self-supervised objectives.

Unlike prior models that treat DDI prediction as a binary classification task [3, 14, 23], our work addresses the overlooked issue of class imbalance across diverse interaction types. Existing approaches often ignore the semantic heterogeneity of DDIs, leading to poor performance on rare but clinically significant classes. We propose a novel integration of GFlowNet and VGAE to explicitly model interaction-type distributions, where GFlowNet generates balanced synthetic samples based on reward-guided sampling. This framework not only improves prediction robustness across all classes but also provides a generalizable strategy for imbalanced biomedical graph problems.

3 Model Architecture

Our framework tackles the class imbalance in Drug-Drug Interaction (DDI) prediction by synergizing a VGAE with a GFlowNet. The VGAE first learns a rich, graph-based latent representation of drugs. The GFlowNet is then trained to generate synthetic DDI samples, focusing on rare yet plausible

interactions. These generated samples augment the original dataset, leading to a more robust and balanced training process for the final prediction model.

3.1 Variational Graph Autoencoder (VGAE) for DDI Representation Learning

We model the set of known DDIs as a multi-relational graph $\mathcal{G} = (\mathcal{D}, \mathcal{E})$, where $\mathcal{D}$ is the set of drug nodes and $\mathcal{E}$ is the set of edges representing interactions. Each edge $(d_i, d_j, t) \in \mathcal{E}$ connects drugs d_i and d_j with a specific interaction type $t \in \mathcal{T}$.

The VGAE architecture consists of a graph-based encoder and a link prediction decoder.

3.1.1 Graph Encoder

The encoder is a Graph Neural Network (GNN), such as a Relational Graph Convolutional Network (R-GCN), that learns a latent representation for each drug. It takes the entire drug graph as input and produces a matrix of latent vectors $\mathbf{Z} \in \mathbb{R}^{|\mathcal{D}| \times K}$, where each row $\mathbf{z}_d$ is the embedding for drug d. The encoder, parameterized by ϕ, defines a variational posterior $q_\phi(\mathbf{Z}|\mathcal{G})$ that approximates the true posterior. We assume a factorized Gaussian distribution:

$$q_\phi(\mathbf{Z}|\mathcal{G}) = \prod_{i=1}^{|\mathcal{D}|} q_\phi(\mathbf{z}_i|\mathcal{G}) = \prod_{i=1}^{|\mathcal{D}|} \mathcal{N}(\mathbf{z}_i|\boldsymbol{\mu}_i, \text{diag}(\boldsymbol{\sigma}_i^2))$$

where the means $\boldsymbol{\mu}$ and variances $\boldsymbol{\sigma}^2$ are the outputs of the GNN.

3.1.2 Link Prediction Decoder

The decoder, parameterized by θ, reconstructs the DDI graph from the latent embeddings. For any pair of drugs (d_i, d_j), it predicts the probability of each interaction type $t \in \mathcal{T}$. We use a multi-relational decoder, such as a DistMult model or a simple Multi-Layer Perceptron (MLP):

$$p_\theta(t|\mathbf{z}_i, \mathbf{z}_j) = \frac{\exp(f_\theta(\mathbf{z}_i, \mathbf{z}_j, t))}{\sum_{t' \in \mathcal{T}} \exp(f_\theta(\mathbf{z}_i, \mathbf{z}_j, t'))}$$

where f_θ is a scoring function (e.g., $\mathbf{z}_i^\top \mathbf{R}_t \mathbf{z}_j$ for DistMult, where $\mathbf{R}_t$ is a diagonal matrix for type t).

3.1.3 VGAE Objective

The VGAE is trained by maximizing the evidence lower bound (ELBO) on the training data $\mathcal{D}_{\text{train}}$:

$$\mathcal{L}_{\text{VGAE}}(\theta, \phi) = \mathbb{E}_{q_\phi(\mathbf{Z}|\mathcal{G}_{\text{train}})} \left[\sum_{(d_i, d_j, t) \in \mathcal{E}_{\text{train}}} \log p_\theta(t|\mathbf{z}_i, \mathbf{z}_j) \right] - \text{KL}\left(q_\phi(\mathbf{Z}|\mathcal{G}_{\text{train}}) \,\|\, p(\mathbf{Z})\right)$$

where $p(\mathbf{Z}) = \prod_i \mathcal{N}(\mathbf{z}_i|0, \mathbf{I})$ is the standard Gaussian prior and KL is the Kullback-Leibler divergence.

3.2 GFlowNet for Balanced Synthetic DDI Generation

The core issue with the initial VGAE is that it will be biased towards frequent interaction types. To mitigate this, we employ a GFlowNet to generate a supplementary dataset of synthetic DDIs, with a preference for rare types.

A GFlowNet learns a policy to construct an object x through a sequence of actions, such that the probability of generating x is proportional to a given reward $R(x)$.

3.2.1 Trajectory and State Space

We define the generation of a DDI triple $x = (d_i, d_j, t)$ as a three-step trajectory:

1. **Initial State (s_0)**: The empty set.
2. **Action 1 (Select Type)**: From s_0, select an interaction type $t \in \mathcal{T}$. The new state is $s_1 = (t)$.

3. **Action 2 (Select First Drug)**: From s_1, select the first drug $d_i \in \mathcal{D}$. The new state is $s_2 = (t, d_i)$.
4. **Action 3 (Select Second Drug)**: From s_2, select the second drug $d_j \in \mathcal{D} \setminus \{d_i\}$. The final state is the terminal state $s_f = (t, d_i, d_j)$.

3.2.2 Reward Function

The reward function $R(x)$ guides the GFlowNet to generate valuable samples. We design a composite reward that balances **rareness** and **plausibility**:

$$R(t, d_i, d_j) = \underbrace{\left(\frac{1}{n_t + 1}\right)^{\alpha}}_{\text{Rareness}} \times \underbrace{p_\theta(t|\mathbf{z}_i, \mathbf{z}_j)}_{\text{Plausibility}}$$

1. **Rareness**: n_t is the frequency of interaction type t in $\mathcal{D}_{\text{train}}$. The hyperparameter $\alpha \geq 0$ controls the strength of the emphasis on rare classes.
2. **Plausibility**: $p_\theta(t|\mathbf{z}_i, \mathbf{z}_j)$ is the probability assigned by the pre-trained VGAE decoder. This ensures that the generated samples are consistent with the learned embedding space.

3.2.3 GFlowNet Policy and Training

The GFlowNet learns a stochastic forward policy $P_F(s'|s; \psi)$, parameterized by a neural network ψ, for transitioning between states. To make training feasible, for Action 3 (selecting d_j), we restrict the action space. Given state $s_2 = (t, d_i)$, the candidate set for d_j is constructed from the K-nearest neighbors of d_i in the VGAE's latent space $\mathbf{Z}$.

We train the policy network ψ using the **Trajectory Balance (TB)** loss. The TB loss enforces a flow-matching condition over complete trajectories $\tau = (s_0 \to s_1 \to \cdots \to s_f)$:

$$\mathcal{L}_{\text{TB}}(\psi) = \left(\log \frac{Z_\psi \prod_{s \to s' \in \tau} P_F(s'|s; \psi)}{R(s_f)}\right)^2$$

Here, Z_ψ is a learnable parameter representing the total flow (partition function). Minimizing this loss encourages the learned sampling distribution $P_{\text{GFN}}(x)$ to be proportional to the reward $R(x)$.

3.3 End-to-End Training Pipeline

The complete training process involves three main stages:

1. **Stage 1: VGAE Pre-training**: The VGAE model (θ, ϕ) is trained on the original, imbalanced dataset $\mathcal{D}_{\text{train}}$ by minimizing $\mathcal{L}_{\text{VGAE}}$. This yields initial drug embeddings $\mathbf{Z}$ and a plausible decoder p_θ.
2. **Stage 2: GFlowNet Training**: Using the pre-trained embeddings $\mathbf{Z}$ and decoder p_θ to compute the reward, the GFlowNet policy $P_F(\cdot|\cdot; \psi)$ is trained by minimizing the Trajectory Balance loss $\mathcal{L}_{\text{TB}}$.
3. **Stage 3: Augmentation and Final VGAE Training**: The trained GFlowNet policy is used to sample a set of N synthetic DDIs, $\mathcal{D}_{\text{synth}}$. The original VGAE model is then re-trained (or fine-tuned) on the augmented dataset $\mathcal{D}_{\text{aug}} = \mathcal{D}_{\text{train}} \cup \mathcal{D}_{\text{synth}}$, again by minimizing $\mathcal{L}_{\text{VGAE}}$. This final model is used for DDI prediction.

3.4 Evaluation Metrics

To evaluate the effectiveness of our proposed generative augmentation strategy in addressing class imbalance in drug–drug interaction (DDI) prediction, we employ two key diversity metrics: *Shannon Entropy (SE)* and *Jensen–Shannon Divergence (JSD)*.

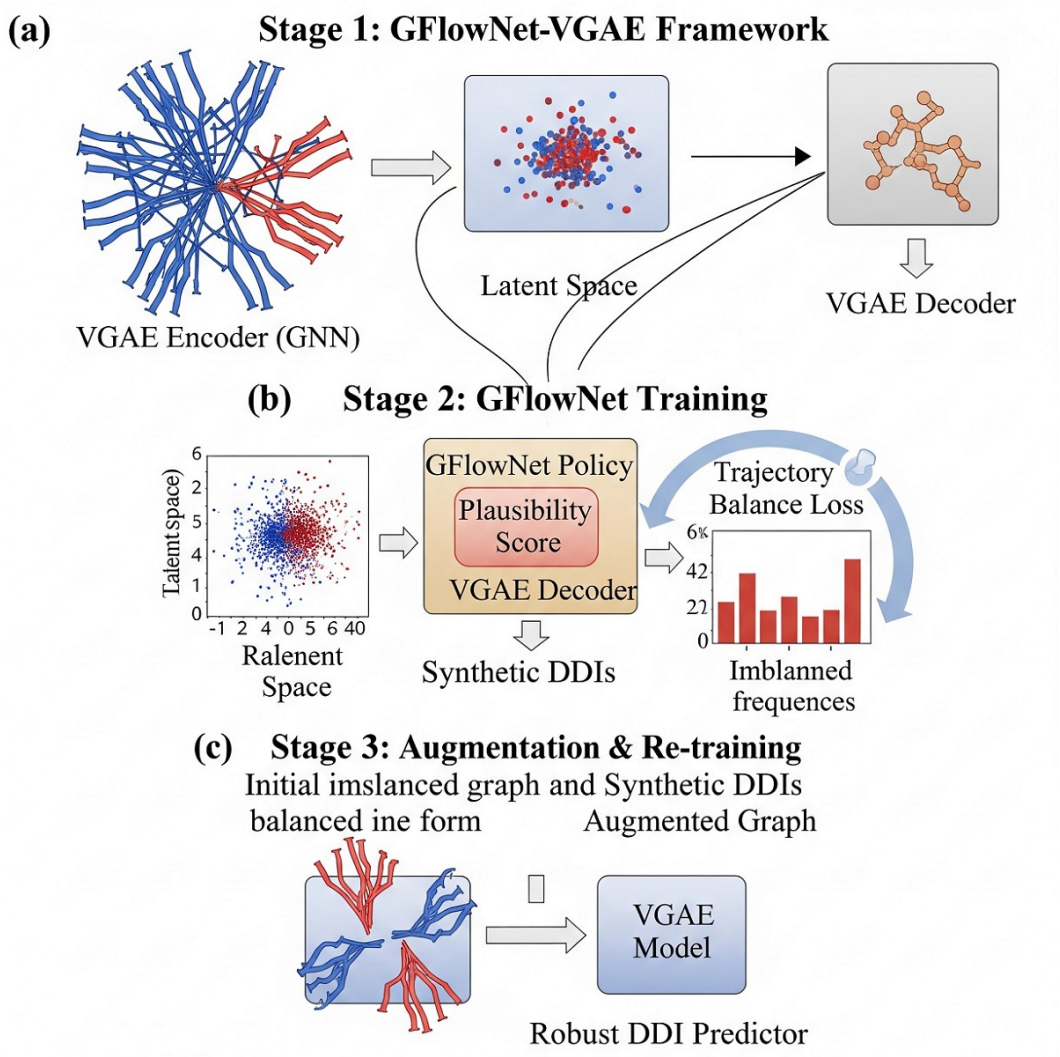

Figure 1: The proposed end-to-end framework. (a) A VGAE is first pre-trained on the original imbalanced DDI graph to learn drug embeddings. (b) A GFlowNet is then trained, using a reward function that combines plausibility from the VGAE and a rareness score, to learn a policy for generating synthetic DDIs. (c) The original data is augmented with the GFlowNet samples and used to train the final, robust VGAE model.

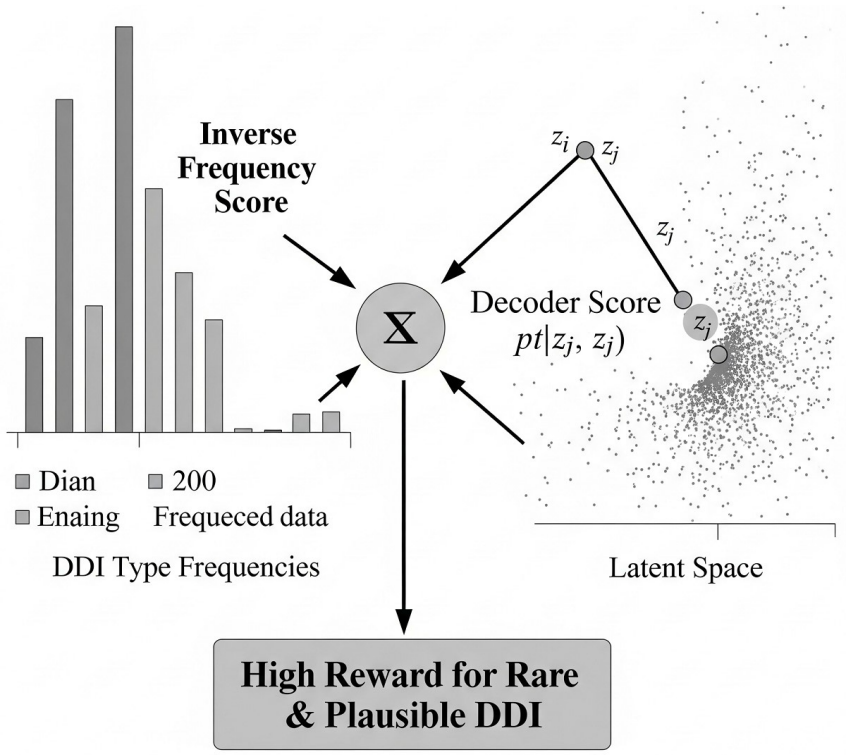

Figure 2: Composition of the GFlowNet reward function. The reward for a generated sample (d_i, d_j, t) is the product of a rareness score, which is inversely proportional to the type's frequency in the training data, and a plausibility score, derived from the pre-trained VGAE decoder's confidence.

3.4.1 Shannon Entropy

Shannon Entropy (SE) [6], measures the uncertainty or diversity of a probability distribution. For a discrete distribution $P = \{p_1, p_2, ..., p_n\}$, the Shannon Entropy is computed as:

$$H(P) = -\sum_{i=1}^{n} p_i \log_2 p_i, \tag{1}$$

where p_i denotes the probability of the i-th interaction type. A higher entropy value indicates a more uniform distribution across interaction types, signifying reduced class imbalance. In our experiments, the SE increased from 1.23 to 1.69 after applying the GFlowNet-based augmentation. This reflects a more balanced representation of interaction types in the dataset, which is especially beneficial for improving the learning of under-represented classes.

3.4.2 Jensen–Shannon Divergence

Jensen–Shannon Divergence (JSD)[13], quantifies the similarity between two probability distributions. For two distributions P (augmented) and Q (true), JSD is defined as:

$$JSD(P||Q) = \frac{1}{2} D_{KL}(P||M) + \frac{1}{2} D_{KL}(Q||M), \tag{2}$$

where $M = \frac{1}{2}(P+Q)$ is the average of the two distributions, and D_{KL} denotes the Kullback–Leibler divergence. JSD is symmetric and bounded between 0 and 1, where lower values indicate greater similarity. In our case, JSD decreased from 0.35 to 0.12, indicating that the distribution of synthetic samples more closely matches the empirical distribution of real interactions. This alignment ensures that the generative process is not just balancing classes arbitrarily but doing so in a way that reflects the true underlying data distribution.

Algorithm 1 End-to-End Training of GFlowNet-VGAE for DDI Prediction

Input: DDI graph $\mathcal{G}_{\text{train}}$, number of synthetic samples N, GFN candidate size K.
Parameters: VGAE parameters (θ, ϕ), GFlowNet policy parameters ψ.
function TRAINVGAE($\mathcal{G}$, epochs)
    Initialize VGAE parameters (θ, ϕ).
    for epoch = 1 to epochs **do**
        Sample a batch of edges from $\mathcal{G}$.
        Compute latent embeddings $\mathbf{Z} \sim q_\phi(\cdot|\mathcal{G})$.
        Calculate $\mathcal{L}_{\text{VGAE}}$ using the batch.
        Update (θ, ϕ) via gradient descent on $\mathcal{L}_{\text{VGAE}}$.
    end for
    return Trained (θ, ϕ) and final embeddings $\mathbf{Z}$.
end function
function TRAINGFLOWNET($\mathcal{G}_{\text{train}}$, (θ, ϕ), $\mathbf{Z}$, epochs)
    Initialize GFlowNet policy parameters ψ and total flow Z_ψ.
    for epoch = 1 to epochs **do**
        // Sample a trajectory τ
        $s_0 \leftarrow \emptyset$
        $t \sim P_F(\cdot|s_0; \psi)$ ▷ Sample interaction type
        $s_1 \leftarrow (t)$
        $d_i \sim P_F(\cdot|s_1; \psi)$ ▷ Sample first drug
        $s_2 \leftarrow (t, d_i)$
        Construct candidate set $\mathcal{C}_{d_i}$ using K-NN of d_i in $\mathbf{Z}$.
        $d_j \sim P_F(\cdot|s_2, \mathcal{C}_{d_i}; \psi)$ ▷ Sample second drug from candidates
        $s_f \leftarrow (t, d_i, d_j)$
        // Compute reward and loss
        Calculate reward $R(s_f)$ using $p_\theta(t|\mathbf{z}_i, \mathbf{z}_j)$ and frequency of t.
        Calculate trajectory log-probability $\log P(\tau) = \sum_{s \to s' \in \tau} \log P_F(s'|s)$.
        Calculate $\mathcal{L}_{\text{TB}} = (\log Z_\psi + \log P(\tau) - \log R(s_f))^2$.
        Update ψ and Z_ψ via gradient descent on $\mathcal{L}_{\text{TB}}$.
    end for
    return Trained policy network ψ.
end function
// Stage 1: Pre-train VGAE
$(\theta_{\text{pre}}, \phi_{\text{pre}}), \mathbf{Z}_{\text{pre}} \leftarrow \text{TrainVGAE}(\mathcal{G}_{\text{train}}, \text{epochs}_1)$
// Stage 2: Train GFlowNet
$\psi_{\text{trained}} \leftarrow \text{TrainGFlowNet}(\mathcal{G}_{\text{train}}, (\theta_{\text{pre}}, \phi_{\text{pre}}), \mathbf{Z}_{\text{pre}}, \text{epochs}_2)$
// Stage 3: Augment and Re-train
$\mathcal{D}_{\text{synth}} \leftarrow \emptyset$
for $i = 1$ to N **do**
    Sample a synthetic DDI (d_i, d_j, t) using the trained policy ψ_{trained}.
    $\mathcal{D}_{\text{synth}} \leftarrow \mathcal{D}_{\text{synth}} \cup \{(d_i, d_j, t)\}$.
end for
$\mathcal{G}_{\text{aug}} \leftarrow \mathcal{G}_{\text{train}} \cup$ edges from $\mathcal{D}_{\text{synth}}$.
$(\theta_{\text{final}}, \phi_{\text{final}}), \mathbf{Z}_{\text{final}} \leftarrow \text{TrainVGAE}(\mathcal{G}_{\text{aug}}, \text{epochs}_3)$ ▷ Re-train on augmented data
Output: The final, robust DDI prediction model $(\theta_{\text{final}}, \phi_{\text{final}})$.

Table 1: Experimental Findings

Setup	AUROC	Accuracy	AUPRC	F1 Score	SE	JSV	Coverage
Without GFN	0.99081	0.96859	0.98861	0.98982	1.23	0.35	0.2441
With GFN	0.99071	0.96792	0.98922	0.99914	1.69	0.12	0.7709

4 Experiments

4.1 Dataset

In this preliminary work, we use the DrugBank dataset [24], which includes 1,703 drugs and 191,870 drug pairs spanning 86 DDI types, along with structural and chemical information. The dataset was split into three subsets: 115,185 drug pairs for training, 38,348 for validation, and 38,337 for testing.

4.2 Implementation Details

We implemented VGAE for DDI prediction [14], given its generative nature. The results in Table 1 show the impact of adding the GFlowNet module to the VGAE framework for DDI prediction. AUROC, Accuracy, AUPRC, and F1 scores were all above 0.99, with minor differences in classification metrics, suggesting GFlowNet had little impact on performance. However, diversity and coverage metrics showed substantial improvements, demonstrating GFlowNet's effectiveness in addressing class imbalance.

5 Findings and Analysis

Table 1 demonstrates the results of our experimentation. These metrics provide quantitative insights into the distributional properties of the interaction types before and after augmentation, beyond what is captured by standard classification metrics.

Together, the increase in Shannon Entropy and the decrease in Jensen–Shannon Divergence demonstrate that our framework successfully promotes a more balanced and representative dataset. This, in turn, facilitates improved generalization, particularly for rare interaction types that are often clinically critical. The observed improvements in coverage—from 0.2441 to 0.7709—further support the model's enhanced capacity to capture and learn from these rare interactions. While traditional metrics showed limited change, these diversity-aware evaluations highlight the strength of generative modeling in addressing real-world challenges in biomedical prediction tasks.

6 Discussion

Our proposed framework addresses a longstanding challenge in drug–drug interaction (DDI) prediction: the pervasive class imbalance that limits model generalization and risks overlooking rare yet clinically significant interactions. By integrating a GFlowNet with a Variational Graph Autoencoder (VGAE), we introduce a principled mechanism for generating synthetic DDI samples conditioned on interaction types. The reward-driven sampling process encourages the generation of under-represented classes, leading to improved balance in the training data, as evidenced by higher Shannon Entropy and lower Jensen–Shannon Divergence. This shift enhances the model's ability to learn from both frequent and rare interactions, which is crucial in clinical contexts where even infrequent interactions can have serious consequences. Our improvements in diversity metrics and coverage suggest that generative methods, when guided by structural priors and task-specific incentives, offer a promising alternative to traditional oversampling or class-weighting strategies.

Beyond immediate performance gains, this work lays the foundation for broader applications in imbalanced biomedical prediction tasks, where generative augmentation can act as a form of informed exploration. By aligning synthetic distributions with real-world priors, our approach supports more equitable representation of minority classes—an issue that extends across genomics, adverse event prediction, and rare disease modeling. Furthermore, the modular design of our framework allows for easy extension: future work could explore integrating domain-specific constraints, leveraging external knowledge graphs, or employing reward shaping to target high-risk interaction types. Ultimately,

our contribution not only advances the technical state-of-the-art but also pushes the field toward more robust and fair AI systems that are better aligned with the needs of healthcare and biomedical decision-making.

7 Conclusion

In this work, we introduce a novel approach to addressing class imbalance in drug–drug interaction (DDI) prediction by integrating a GFlowNet-based generative sampling strategy into a variational graph autoencoder (VGAE) framework. Unlike conventional methods that often overfit to dominant classes, our approach leverages the exploration capabilities of GFlowNets to more effectively generate samples representing under-represented interaction types. While standard classification metrics, such as accuracy and AUC, remain largely unchanged, diversity-aware evaluations reveal substantial gains, underscoring GFlowNet's strength in uncovering rare but clinically significant interactions. This improvement is particularly important for patient safety, as rare DDIs can pose severe risks in real-world clinical settings if left undetected. Moreover, the integration of GFlowNets into the VGAE pipeline demonstrates the potential of generative sampling not only to enhance predictive robustness but also to broaden the representational coverage of interaction spaces.

References

[1] Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J. Hu, Mo Tiwari, and Emmanuel Bengio. Gflownet foundations. *Journal of Machine Learning Research*, 24(210):1–55, 2023.

[2] Jiatao Chen, Liang Zhang, Ke Cheng, Bo Jin, Xinjiang Lu, and Chao Che. Predicting drug-target interaction via self-supervised learning. *IEEE/ACM transactions on computational biology and bioinformatics*, PP, 2022.

[3] Yuanfei Dai, Chenhao Guo, Wenzhong Guo, and Carsten Eickhoff. Drug–drug interaction prediction with Wasserstein Adversarial Autoencoder-based knowledge graph embeddings. *Briefings in Bioinformatics*, 22(4):bbaa256, 2020.

[4] David Duvenaud, Dougal Maclaurin, Jorge Aguilera-Iparraguirre, Rafael Gómez-Bombarelli, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P. Adams. Convolutional networks on graphs for learning molecular fingerprints. In *Proceedings of the 29th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'15, page 2224–2232, Cambridge, MA, USA, 2015. MIT Press.

[5] Ali Ezzat, Min Wu, Xiao-Li Li, and Chee-Keong Kwoh. Drug-target interaction prediction via class imbalance-aware ensemble learning. *BMC Bioinformatics*, 17(S19), December 2016.

[6] Shu-Cherng Fang and Jacob H. S. Tsao. Entropy optimization: Shannon measure of entropy and its properties. In *Encyclopedia of Optimization*, page 916–921. Springer US, 2008.

[7] Zhenyu Hou, Xiao Liu, Yukuo Cen, Yuxiao Dong, Hongxia Yang, Chunjie Wang, and Jie Tang. Graphmae: Self-supervised masked graph autoencoders. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '22, page 594–604, New York, NY, USA, 2022. Association for Computing Machinery.

[8] Moksh Jain, Tristan Deleu, Jason Hartford, Cheng-Hao Liu, Alex Hernandez-Garcia, and Yoshua Bengio. Gflownets for ai-driven scientific discovery. *Digital Discovery*, 2(3):557–577, 2023.

[9] Dongkwan Kim and Alice Oh. How to find your friendly neighborhood: Graph attention design with self-supervision. In *International Conference on Learning Representations*, 2021.

[10] Thomas N. Kipf and Max Welling. Variational graph auto-encoders, 2016.

[11] Eli Laird, Ayesh Madushanka, Elfi Kraka, and Corey Clark. Xinsight: Revealing model insights for gnns with flow-based explanations, 2023.

[12] Jaechang Lim, Seongok Ryu, Kyubyong Park, Yo Joong Choe, Jiyeon Ham, and Woo Youn Kim. Predicting drug-target interaction using a novel graph neural network with 3d structure-embedded graph representation. *Journal of chemical information and modeling*, 2019.

[13] M.L. Menéndez, J.A. Pardo, L. Pardo, and M.C. Pardo. The jensen-shannon divergence. *Journal of the Franklin Institute*, 334(2):307–318, March 1997.

[14] Nhat Khang Ngo, Truong Son Hy, and Risi Kondor. Predicting drug-drug interactions using deep generative models on graphs, 2022.

[15] Thin Nguyen, Hang Le, Thomas P Quinn, Tri Nguyen, Thuc Duy Le, and Svetha Venkatesh. GraphDTA: predicting drug–target binding affinity with graph neural networks. *Bioinformatics*, 37(8):1140–1147, 2020.

[16] Andrei Cristian Nica, Moksh Jain, Emmanuel Bengio, Cheng-Hao Liu, Maksym Korablyov, Michael M. Bronstein, and Yoshua Bengio. Evaluating generalization in GFlownets for molecule design. In *ICLR2022 Machine Learning for Drug Discovery*, 2022.

[17] Mohit Pandey, Gopeshh Subbaraj, and Emmanuel Bengio. Gflownet pretraining with inexpensive rewards, 2024.

[18] Julien Roy, Pierre-Luc Bacon, Christopher Pal, and Emmanuel Bengio. Goal-conditioned gflownets for controllable multi-objective molecular design, 2023.

[19] Seonghwan Seo, Minsu Kim, Tony Shen, Martin Ester, Jinkyoo Park, Sungsoo Ahn, and Woo Youn Kim. Generative flows on synthetic pathway for drug design. In *The Thirteenth International Conference on Learning Representations*, 2025.

[20] Martin Simonovsky and Nikos Komodakis. Graphvae: Towards generation of small graphs using variational autoencoders. In *International Conference on Artificial Neural Networks*, 2018.

[21] Wen Torng and Russ B. Altman. Graph convolutional neural networks for predicting drug-target interactions. *bioRxiv*, 2018.

[22] Xiao Wang, Nian Liu, Hui jun Han, and Chuan Shi. Self-supervised heterogeneous graph neural network with co-contrastive learning. *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2021.

[23] Azmine Toushik Wasi, Taki Hasan Rafi, Raima Islam, Serbetar Karlo, and Dong-Kyu Chae. Cadgl: Context-aware deep graph learning for predicting drug-drug interactions, 2024.

[24] David S Wishart, Yannick D Feunang, An C Guo, Elvis J Lo, Ana Marcu, Jason R Grant, Tanvir Sajed, Daniel Johnson, Carin Li, Zinat Sayeeda, et al. Drugbank 5.0: a major update to the drugbank database for 2018. *Nucleic acids research*, 46(D1):D1074–D1082, 2018.

[25] Yaming Yang, Ziyu Guan, Zhe Wang, Wei Zhao, Cai Xu, Weigang Lu, and Jianbin Huang. Self-supervised heterogeneous graph pre-training based on structural clustering. *Advances in Neural Information Processing Systems*, 35:16962–16974, 2022.

[26] Ziduo Yang, Weihe Zhong, Qiujie Lv, and Calvin Yu-Chian Chen. Learning size-adaptive molecular substructures for explainable drug–drug interaction prediction by substructure-aware graph neural network. *Chemical Science*, 13(29):8693–8703, 2022.

[27] Yan Zhao, Jun Yin, Li Zhang, Yong Zhang, and Xing Chen. Drug–drug interaction prediction: databases, web servers and computational models. *Briefings in Bioinformatics*, 25(1):bbad445, 2023.